

AI Policy Development Guide

A starting point for child and family service organizations at any stage of the AI journey.

Staff are already using AI tools today, with or without guidance. The question isn't whether your organization will engage with AI, it's whether that engagement is deliberate. A written policy turns "we're figuring it out" into a clear, documented stance your board, funders, and MCP partners can see.

The 3-Stage Framework

- 1 Understand It Yourself**
Build a personal baseline: try a real nonprofit-relevant prompt yourself before bringing anything to your team.
- 2 Establish Your Guardrails**
A policy doesn't need to be long, it needs to answer five questions clearly enough that any staff member can act on it.
- 3 Identify Your First Use Case**
Start at the low-risk end of the ladder below: no client PHI, a real recurring burden, and a human review before anything ships.

First Use Cases — In Order of Risk

Start at 1. Earn your way to 5.

- 1 Grant research & drafting in an enterprise chatbot**
No client data; human review before anything ships.
- 2 Staff training content**
Internal-facing; mistakes get caught in review.
- 3 Family FAQs in plain language**
External-facing; review for accuracy and reading level.
- 4 Meeting summaries**
Approved tool only, with your consent process in use.
- 5 Clinical documentation & case notes**
The biggest prize, and last for a reason. PHI at the center of the legal record: compliant vendor, BAA, consent workflow, and a pilot first.

5 Components Every Policy Should Cover

Approved Tools

What staff may and may not use.

Prohibited Uses

No client PHI in consumer tools; no unapproved automation.

Confidentiality

How HIPAA, CMIA, and 42 CFR Part 2 apply to AI-assisted work.

Staff Training

Expectations and accountability for anyone using an approved tool.

Review Cadence

AI tools change rapidly; set a date to revisit the policy.

Building Your Policy: Eight Decisions

Write these eight decisions down and you have a policy your board can approve. A useful policy is short. Fill in the eight items, keep the tool list current, and you have something staff can follow and a board can sign.

These eight decisions map to the AI standard CARF added to its ASPIRE to Excellence framework, effective July 1, 2026. If your organization is CARF-accredited and uses AI in service delivery, written policies and procedures are now required.

1 Sort your information into levels

A few clear categories, from regulated client data down to public material. The rules follow the information, not the device it is on.

2 List the tools staff may use

If it is not on the list, it is not approved. Confirm each vendor agreement actually covers the AI features, not just the product.

3 Choose one path for sensitive information

Name a single approved tool for your most protected data. One clear path is followed more reliably than a set of exceptions.

4 State what is off limits

Short, absolute rules that need no interpretation. A starting set is below.

5 Say what happens when something goes wrong

Stop, do not delete anything, report right away. Name a contact and a number, and promise no retaliation for honest reporting.

6 Name an owner and a governance committee

One accountable person, plus a small committee that reviews vendors and data practices. The owner convenes it. Review policies at least annually, and again after any new tool, rule change, or close call.

7 Tell the people you serve

Disclose your use of AI in service delivery, and separately for any analysis or processing of their health information. Recording consent is not the whole obligation. Write the disclosure language once and use it consistently.

8 Train the users and the reviewers

Regular training for staff using AI tools, and separate training for the people responsible for oversight. Document who was trained and when.

Your Tool List

Tool	Vendor	Approved For	Who May Use It

A Starting Set of Off-Limits Rules

- No client or organizational information in personal or consumer AI accounts.
- No AI-generated clinical judgment — no diagnoses, assessments, or treatment decisions.
- Nothing AI produces gets filed or sent until a qualified person has reviewed and edited it. Someone is also accountable for reviewing AI output quality on an ongoing basis, not just item by item.
- No AI note-takers or meeting assistants in client sessions until a documented consent process is in place.
- Your most legally protected records stay out of every AI tool, without exception.

Before You Publish

- Claim only the safeguards you have actually turned on. A policy describing controls you do not have is a liability, not a protection.
- If you are accredited, expect to be asked during survey how your AI use is governed. A policy you cannot demonstrate in practice is worse than no policy.
- Give staff a procedure for anything you require of them, not just the requirement.
- Confirm what your vendor agreements cover. Do not assume.

A starting point, not legal advice. This guide is an educational resource. Your obligations depend on your programs, contracts, funders, and the populations you serve. Review any policy you adopt with your compliance officer and legal counsel before publishing — and revisit it with them as tools and regulations change. This guide is aligned to the CARF AI standard but is not a substitute for the standards manual or for guidance from your accreditation consultant.

FREE FOR ALLIANCE MEMBERS

30-Minute AI Consultation for Alliance Members

Schedule yours: we'll help turn this guide into a policy your board can sign off on.

Your current AI usage · Policy gaps · Board-ready next steps

[Schedule Your Consultation](#)

Steven Hart, President · shart@advantagemicro.net



Scan to schedule